Georgia Tech

Advanced Performance Modeling

Alex Sim, W. Wiliam Yoo, **SDM, CRD, LBNL**
Constantine Dovrolis, Danny Lee, Kamal Shadi, **Georgia Tech**

apm@hpcrd.lbl.gov

Advanced Performance Modeling

Project Goals

- Develop performance prediction models and software framework for high-bandwidth networks
- Develop a performance prediction tool for network traffic

Edge-to-Edge large transfer identification and throughput inference based on link utilization counts

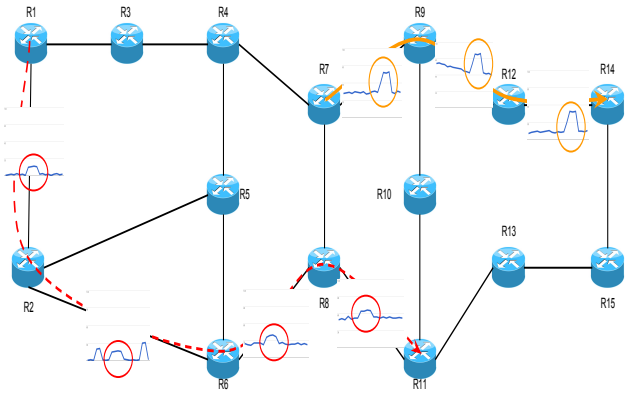

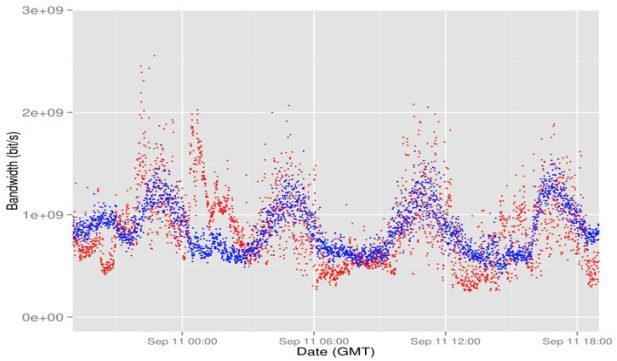

Network traffic prediction model, validated with the actual traffic. Error is within the variance. Blue is the prediction and red is observed traffic

Accomplishments

- Developed overall performance inference and prediction framework
- Inference of edge-to-edge network traffic
  - * Enabling prediction, tracing and quantifying the network traffic with partial observations
  - * Edge-to-edge transfer throughput inference using link utilization counts
- Network traffic prediction model
  - * Improved efficiency and accuracy of prediction by adaptive models
- Over the last ~3 years, 7 paper publications, 4 papers in preparation, 2 invited talks, 1 provisional patent, 2 software prototypes

Impact

- Enable scientific collaborations to utilize the resources offered by high-bandwidth network infrastructures more effectively
  - * Improve network usage and enable predictable data throughput
  - * Long-term capacity and traffic engineering planning of network infrastructures

Prototypes

Advanced Performance Model (APM)

Bandwidth Utilization Prediction Prototype

To improve the efficiency of resource utilization and scheduling of scientific data transfers on high-speed networks, we started a project on Advanced Performance Modeling (APM) that investigates and models a general-purpose network performance estimation framework. The performance prediction model and the framework will be helpful in optimizing the performance and utilization of fast networks as well as sharing resources with predictable performance for scientific collaborations, especially in data intensive applications. The prediction model estimates future network usage and the latency in using the network. Historical network performance information is used for throughput prediction without putting extra load on the resources by active measurement collection.
The details are describe in "Network Bandwidth Utilization Forecast Model on High Bandwidth Network", LBNL-6677E, 2014. (PDF)

ANL -> ORNL

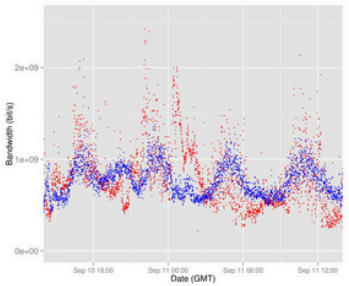

Blue is the prediction of one day ahead of the current time.
Red is the observed data of the same duration.

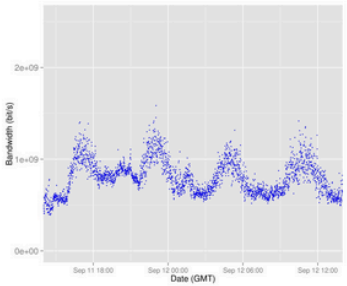

The prediction of one day ahead of the current time.

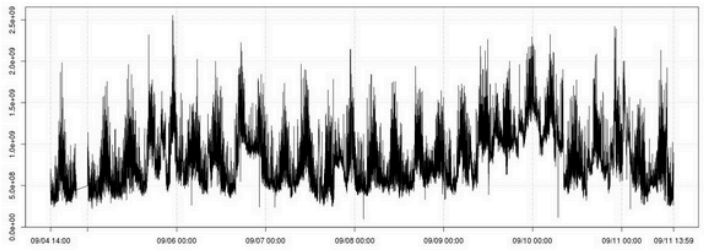

The bandwidth Utilization during 1 week before prediction

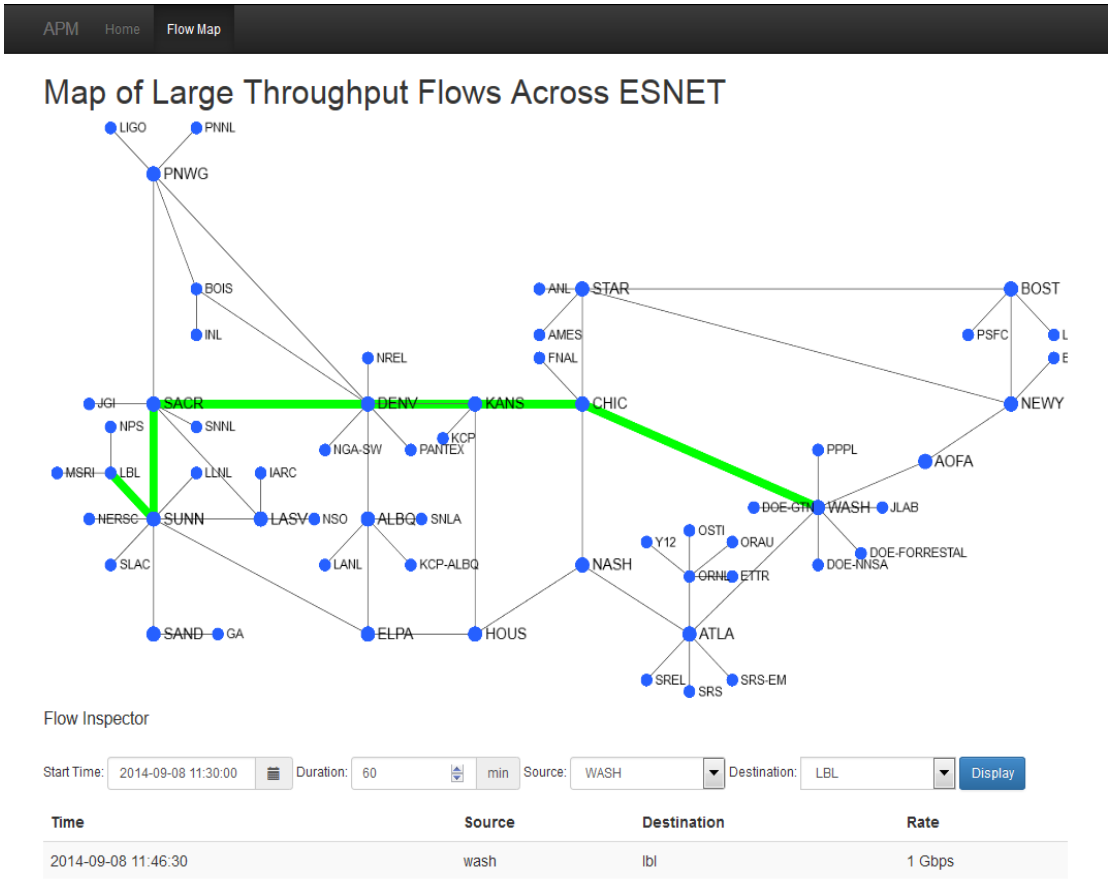

Edge-to-Edge Transfer Throughput Inference Using Link Utilization Counts

The APM Team at Georgia Tech

Danny Lee

Kamal Shadi

Kanu Sahai

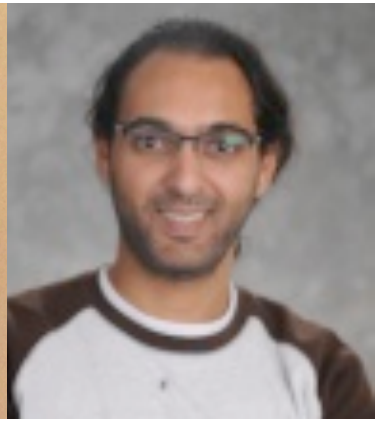

Demetris Antoniades
(former)

Warren Matthews
(former)

What are we trying to do?

- The Google Maps metaphor
  - Real-time view of traffic in road network
  - Average speed of vehicles in each road segment (not just speed limit)
  - Travel-time prediction for given start-destination points
  - Relies on information from moving vehicles that carry a smart phone
  - No "active probing"

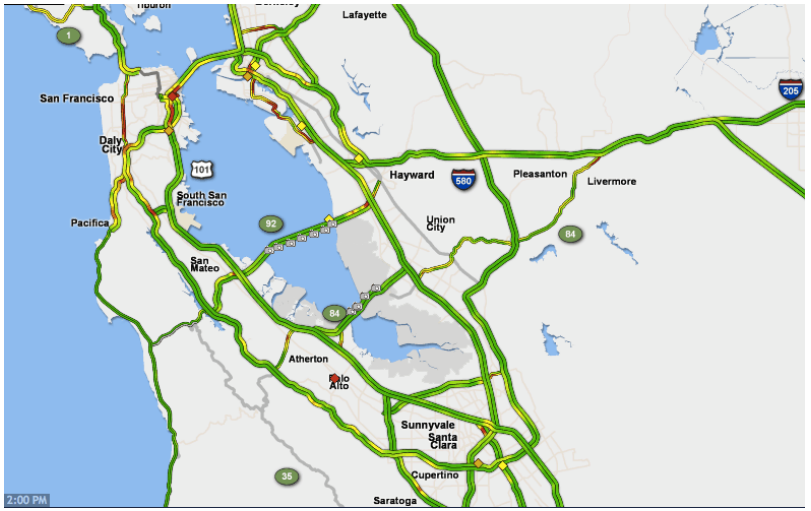

- How can we do the same in a computer network (say ESnet?)
  - Active probing? Slow & intrusive
  - Netflow-based samples? Privacy and data management issues
  - Netflow focuses on flows, not multi-connection transfers
  - Is there a better/simpler way?

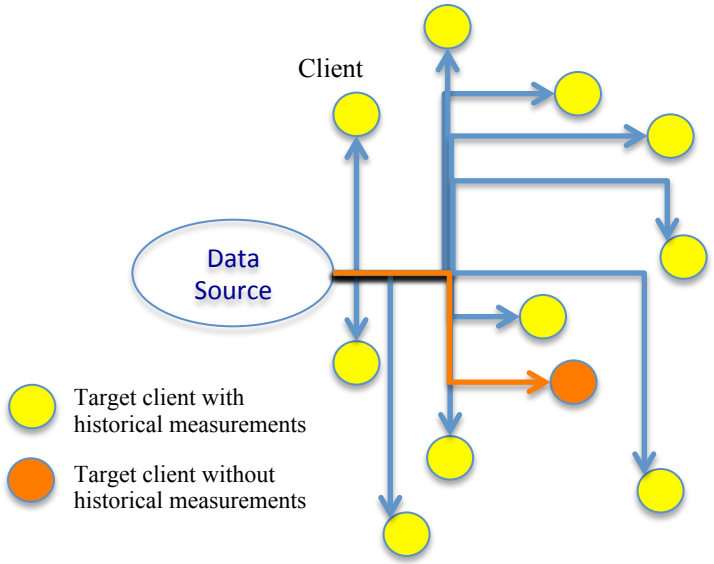

Objective

- Create a method that uses per-link aggregate utilization data to:
  - Identify large transfers and estimate their achieved throughput by observing variations in the aggregated throughput
  - Track these transfers through the network and identify ingress and egress interfaces for each large transfer
  - Rely only on SNMP-based link utilization data (low-cost, no privacy issues)

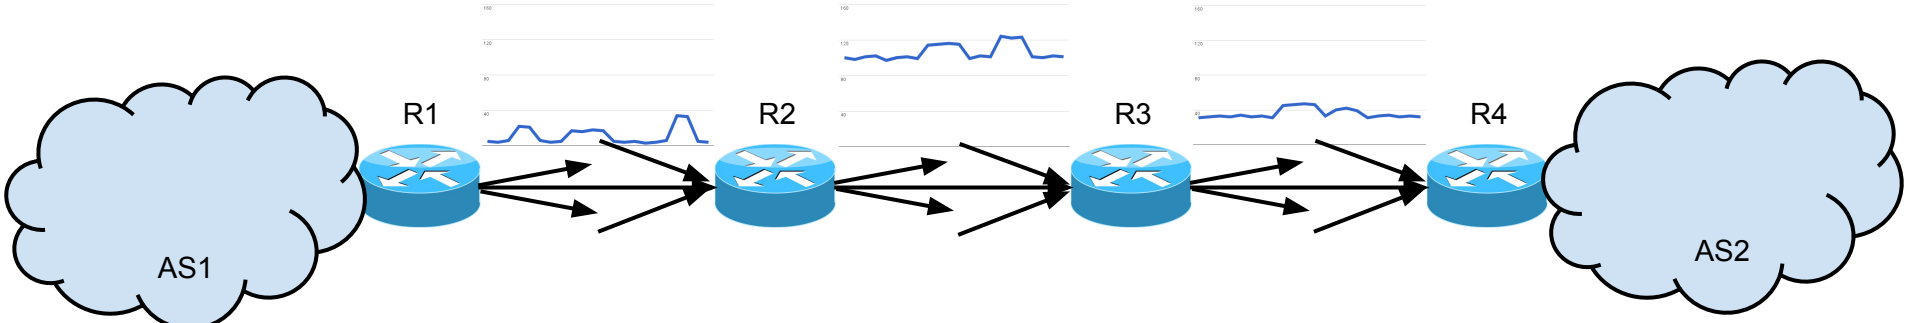

Two basic ideas behind this method

- Major & sudden deviations in link aggregate utilization correspond to beginning/end of major network transfers
  - Not just individual flows
- The path of a large transfer can be identified:
  - match a major load deviation at an incoming link with a similar major deviation at an outgoing link of the same router at about the same time

Illustration

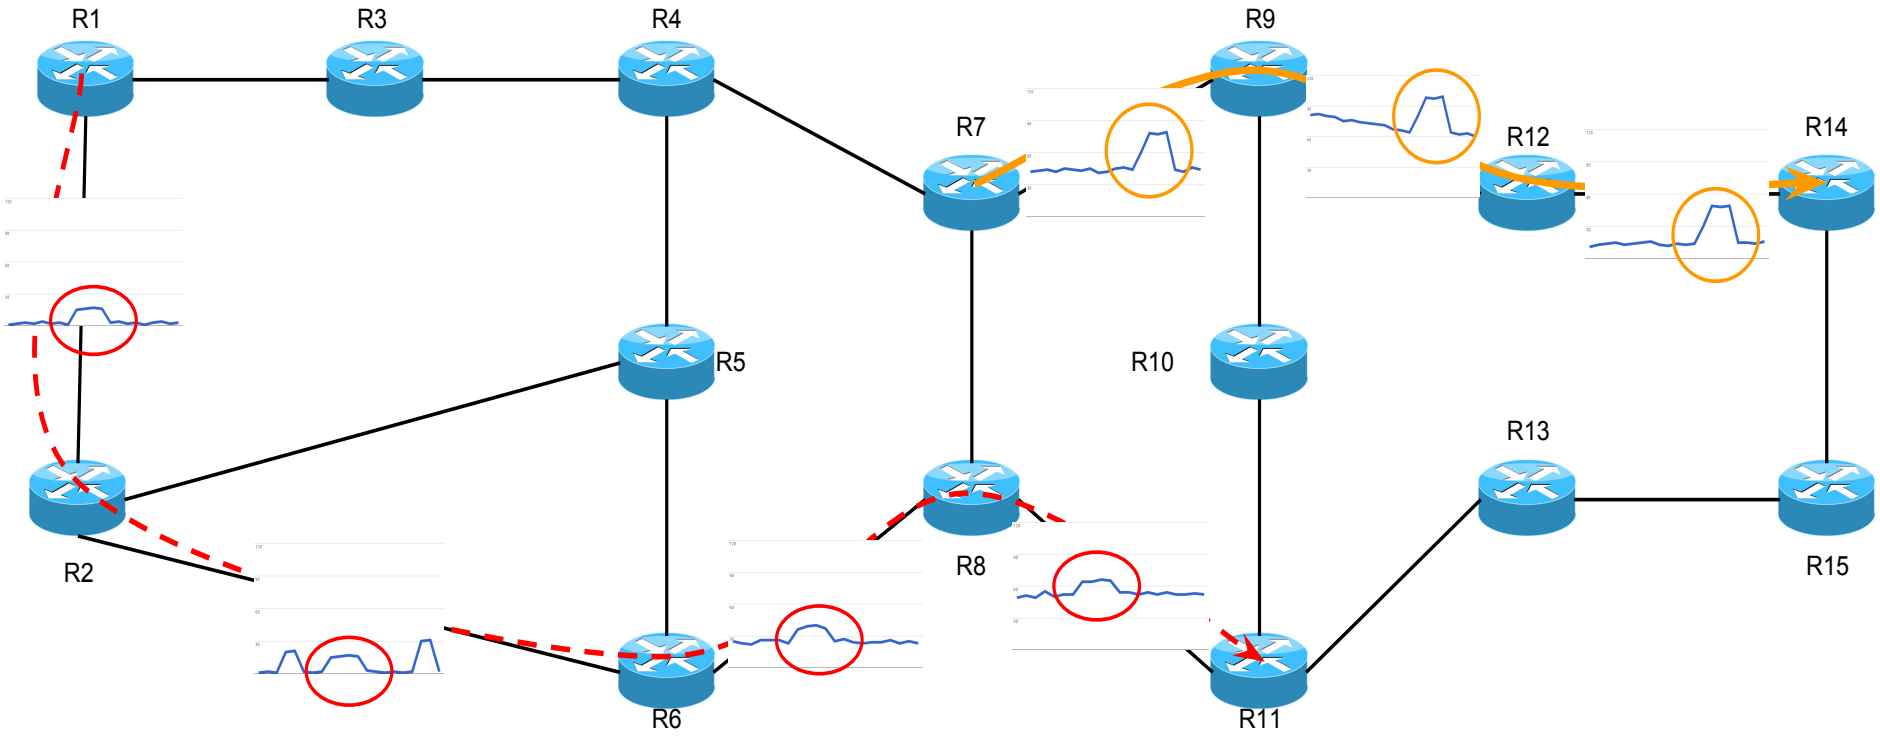

See poster for more details

1. A statistical methodology to identify major events in per-link utilization time series
2. An algorithm to map events in an input interface of a router to the output interfaces of the same router the event is switched to
3. Evaluation based on ESnet data and experiments shows that method can accurately identify transfers that last > 2minutes & achieve > 3Mbps.

Some Applications

- Edge-to-edge throughput prediction
  - Use identified transfers as samples for throughput prediction methods
  - User specifies ingress-egress paid & time period
- Create real-time map of network performance based on ACTUAL ACHIEVED transfer throughputs
- DDoS attack initiator inference
  - Follow identified transfers to the ingress link(s)
  - Identify actual source of spoofed DDOS traffic

Current status

- We are building a system that allows user to:
  - specify ingress-egress interfaces at ESnet map and get throughput prediction for that path based on recent large transfers
  - or, that visualizes e2e throughput levels for all egress links, given a single ingress link

- Based on publicly available SNMP-based link utilization data from ESnet

What question does your research motivate you to now ask?

- Most of (at least) our research so far has been hypothesis-driven
- But we increasingly have access to rich multimodal spatio-temporal network data:
  - PerfSONAR measurements, SNMP-based per-link data, router events, netflow-logs, etc
- What would we be able to learn about a network or distributed application if we could take a hypothesis-free approach, powered by the same tools that drive the "big data revolution"?

Network Traffic Prediction Model

W. William Yoo
SDM, CRD, LBNL

Former team members
Jaesik Choi at UNIST
Kejia Hu at Northwestern Univ.

Motivation

- **Increasing Data Volume**
  - **Efficient resource management and scheduling data movement**
    - **Predict the network bandwidth utilization between two HPC sites**
- **Challenge**
  - **Accurate and fine-grained performance model**
    - **Computational complexities and variances/burstiness**

Bandwidth Utilization

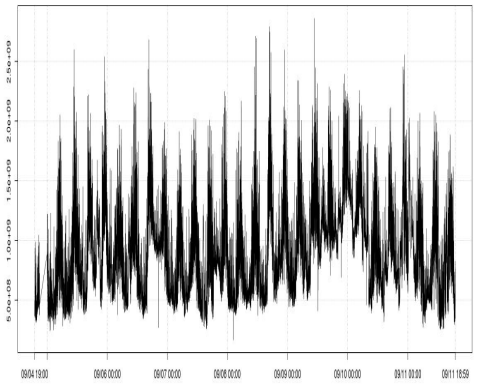

NERSC ➔ ANL

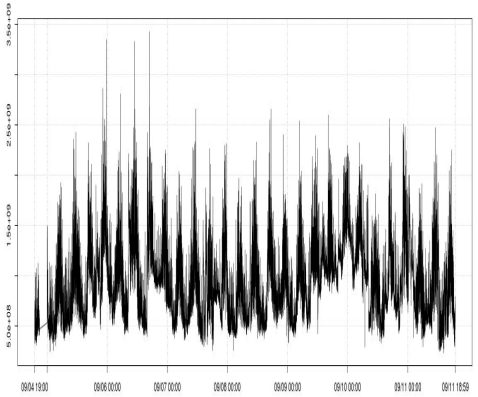

ANL ➔ NERSC

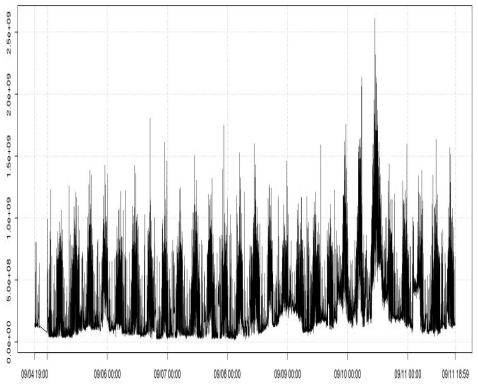

NERSC ➔ ORNL

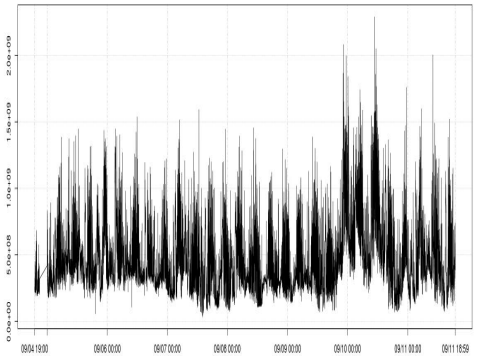

ORNL ➔ NERSC

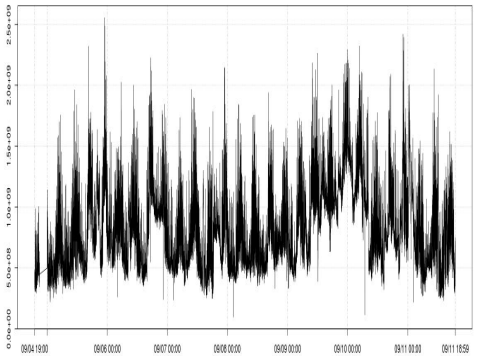

ANL ➔ ORNL

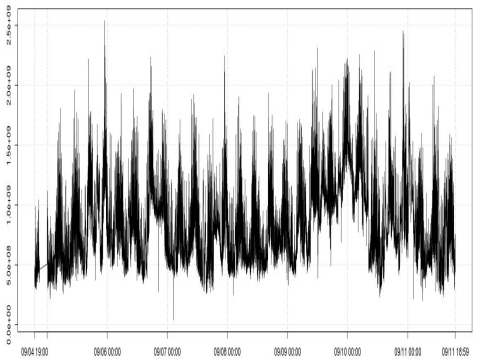

ORNL ➔ ANL

Network Traffic Performance Model

- **Time series model**
  - **ARIMA with STL, logit transformation, and stationarity**
- **Evaluation**

PID	Source	Destination	# of Links
P1	$NERSC$	ANL	7
P2	ANL	$NERSC$	7
P3	$NERSC$	$ORNL$	7
P4	$ORNL$	$NERSC$	7
P5	ANL	$ORNL$	6
P6	$ORNL$	ANL	6

Model Evaluation

Model Comparison

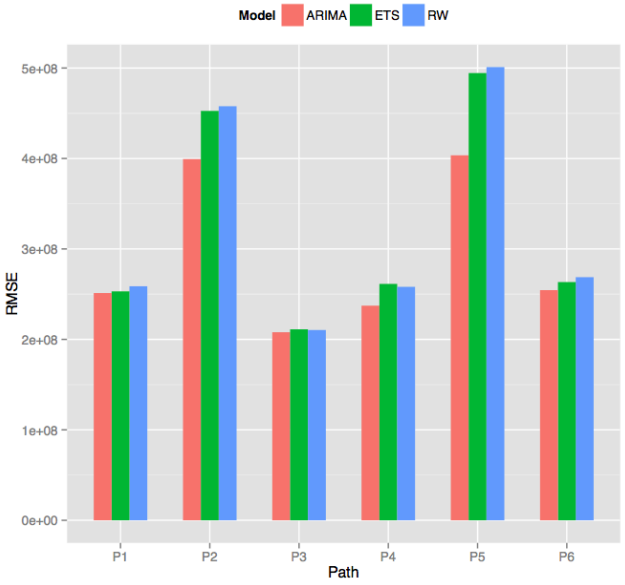

Logit Transformation

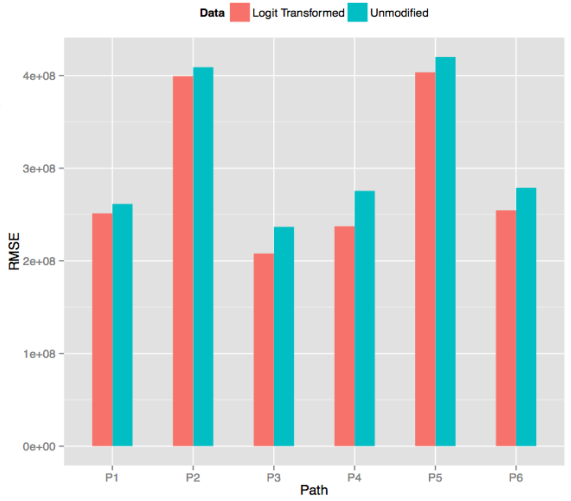

Stationarity

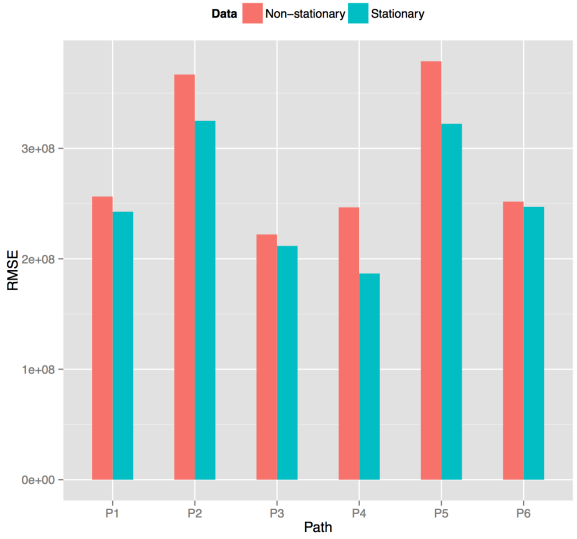

Results

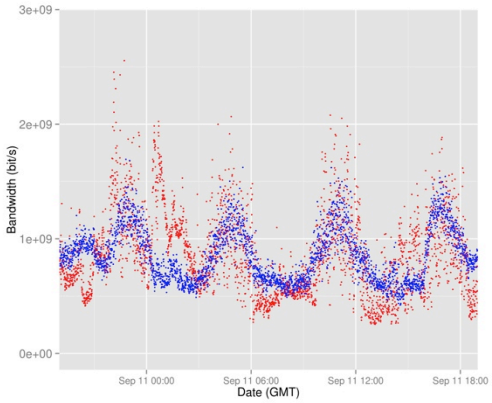

NERSC ➔ ANL

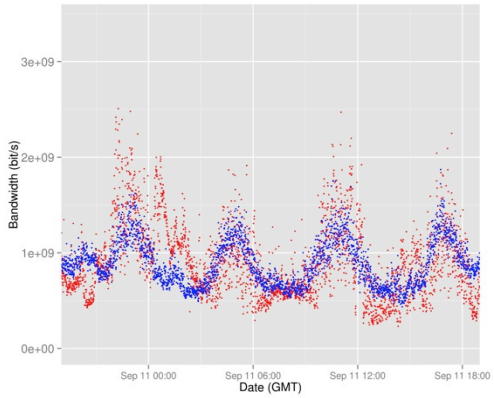

ANL ➔ NERSC

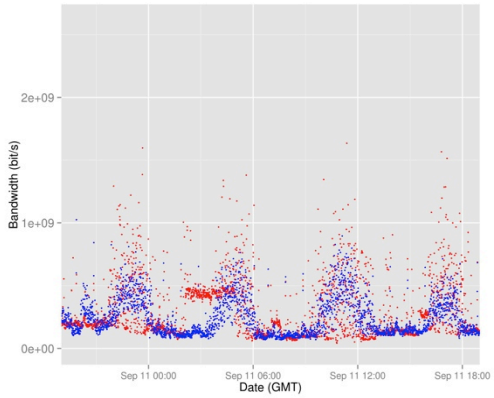

NERSC ➔ ORNL

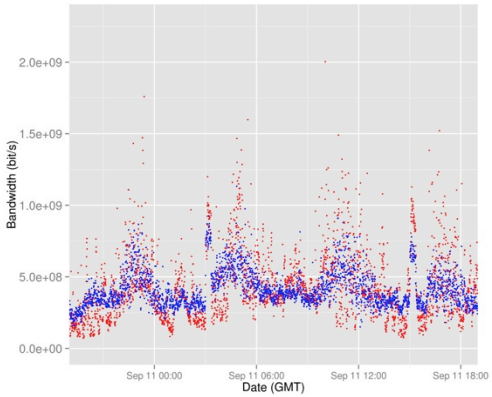

ORNL ➔ NERSC

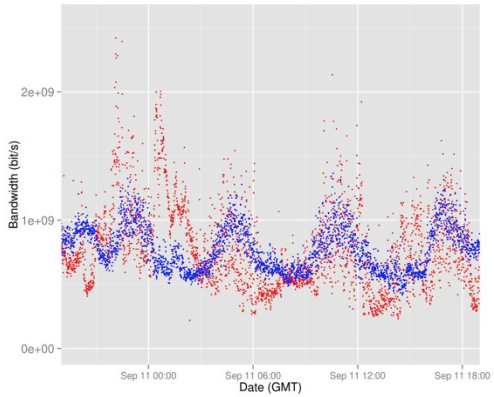

ANL ➔ ORNL

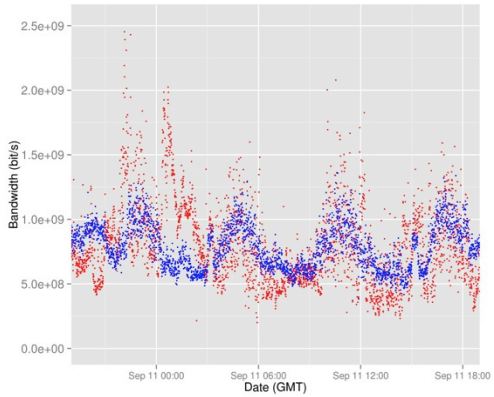

ORNL ➔ ANL

Summary

- **Performance Prediction Model**
  - ARIMA with STL, logit transformation, and stationarity
  - Prediction error is within the variances of observed data
  - Logit transform reduced prediction error by 8.5%
  - Stationarity assumption reduced prediction error by 10.9%
- **Ongoing Work**
  - Adaptive model
    - To adapt the long-term trend changes
  - Change point detection
    - To detect behavioral changes
  - Multivariate prediction model
    - To study correlations between data collections
- **Demo**

Next Research Question

- **Delay to access measurement data**
  - **What are the strategies for managing data repository for efficient access?**
- **Not enough to model for end-to-end data flow performance**
  - **Measurement collections and accessibility**